Measuring Relationship Anonymity in Mix Networks

Vitaly Shmatikov and Ming-Hsiu Wang

The University of Texas at Austin

ABSTRACT

Many applications of mix networks such as anonymous Web browsing require *relationship anonymity*: it should be hard for the attacker to determine who is communicating with whom. Conventional methods for measuring anonymity, however, focus on *sender anonymity* instead. Sender anonymity guarantees that it is difficult for the attacker to determine the origin of any given message exiting the mix network, but this may not be sufficient to ensure relationship anonymity. Even if the attacker cannot identify the origin of messages arriving to some destination, relationship anonymity will fail if he can determine with high probability that at least one of the messages originated from a particular sender, without necessarily being able to recognize this message among others.

We give a formal definition and a calculation methodology for relationship anonymity. Our techniques are similar to those used for sender anonymity, but, unlike sender anonymity, relationship anonymity is sensitive to the distribution of message destinations. In particular, Zipfian distributions with skew values characteristic of Web browsing provide especially poor relationship anonymity. Our methodology takes route selection algorithms into account, and incorporates information-theoretic metrics such as entropy and min-entropy. We illustrate our methodology by calculating relationship anonymity in several simulated mix networks.

Categories and Subject Descriptors: C.2.0[**Computer-Communication Networks**]: Security and protection; K.6.5[**Security and Protection**]

General Terms: Security

Keywords: Anonymity, Privacy, Mix Networks

1. INTRODUCTION

Mix networks, first proposed by Chaum [3], are a practical way of achieving anonymity on insecure communication networks. Intuitively, a *mix* is a server that accepts several incoming messages and forwards them to their respective destinations in such a way that an outside observer cannot link an outgoing message with an incoming message. Mixes are typically assembled into networks, intended to provide some degree of anonymity to their users even if some of the mixes are controlled by the adversary.

In this paper, we focus on *relationship anonymity*. As defined by Pfitzmann *et al.* [7], "relationship anonymity means that it is untraceable who communicates with whom." This is an important property for many practical applications of mix networks. For example, users of an anonymous Web browsing or email system often wish to hide the fact that they are communicating with a particular destination. Other definitions of anonymity address a similar, but slightly different property. For example, Serjantov and Danezis [9] consider the attacker's probability distribution over all possible senders and recipients of a *given message*. Whereas relationship anonymity hides the fact that party A is communicating with party B, anonymity of an individual message hides the fact that A sent it (in the case of sender anonymity) or that B is its intended destination (in the case of recipient anonymity).

Sender anonymity and relationship anonymity are not directly comparable. Consider a set of senders from an oppressive country who are all anonymously accessing a single politically sensitive website. Suppose the network provides perfect sender anonymity, *i.e.*, any message exiting the network is equally likely to have originated from any active sender. By observing these messages, however, the attacker can easily infer that all of them have the same destination. For every active sender, the attacker can thus determine with 100% certainty that this sender is communicating with the website, completely breaking relationship anonymity.

This artificial example indicates that (a) sender anonymity is not sufficient for either recipient anonymity, or relationship anonymity; (b) unlike sender anonymity, relationship anonymity is sensitive to the distribution of potential message destinations; and (c) under some destination distributions, even a perfectly secure, "black-box" mix network cannot guarantee relationship anonymity.

We are also interested in the property called *beyond suspicion* in [8]: the destination with which the user is communicating should not appear significantly more likely than any other possible destination. (We will use an even stronger property that *no* destination should appear more likely than others.) Standard metrics such as entropy of the attacker's a-posteriori distribution of potential destinations do not capture this property. For example, it is entirely possible that in a high-entropy distribution some destination is associated with a probability which is an order of magnitude higher than the probability of any other destination. This calls for alternative anonymity metrics that better capture the ratio between probabilities associated with different members of the anonymity set.

Our techniques follow the basic framework of [9], with the additional emphasis on min-entropy as a measure of anonymity. The use of min-entropy was previously proposed by Tóth *et al.* [10], but in a very different mix network model.

WPES'06, October 30, 2006, Alexandria, Virginia, USA.
Copyright 2006 ACM 1-59593-556-8/06/0010 ...$5.00.

2. DEFINITIONS AND METRICS

Consider a mix network with N senders $S_1, \ldots, S_N$ and H destinations $D_1, \ldots, D_H$. The natural way to define relationship anonymity is via the attacker's a-posteriori probability RA_{ij} (*i.e.*, probability measured after the attacker has completed his observations of the mix network) that the ith sender is communicating with jth destination, where $1 \leq i \leq N, 1 \leq j \leq H$. Because different senders may send a different number of messages, $[\mathsf{RA}_{i1} \ldots \mathsf{RA}_{iH}]$ is not a proper probability distribution: for each sender i, RA_{ij} values add up to the number of messages sent by that sender.

Assume that each sender S_i sends n_i messages $x_{i1}, \ldots, x_{in_i}$. Each message x_{ik} has one destination. Let $\mathsf{RAmsg}_{ik}[1..H]$ be the probability distribution of its potential destinations.

Entropy vs. min-entropy. Given a message and the probability distribution of potential destinations of this message, the standard information-theoretic measure of anonymity is *entropy* of this distribution [9, 4]. Informally, entropy is a measure of how "random" the distribution is. A high entropy value implies that the network provides a high level of anonymity. Entropy of the distribution $\mathsf{RAmsg}_x[1..H]$, where $\mathsf{RAmsg}_x[j]$ is the probability that the destination of some message x is D_j, is calculated as

$$\mathsf{RAent}_x = -\sum_{1 \leq j \leq H} \mathsf{RAmsg}_x[j] \log_2(\mathsf{RAmsg}_x[j])$$

An intuitive interpretation of entropy is that it represents the logarithm of the effective size of the *anonymity set* for the sender of a message. For example, a distribution of potential senders whose entropy is 6 can be interpreted as saying that the sender is indistinguishable from 63 $(= 2^6 - 1)$ other senders.

Entropy does not always capture the right anonymity property. Consider a distribution of 100 potential destinations, in which all but one are equally likely with probability 0.009, and a single destination has probability 0.109. Entropy of this distribution is 6.40, close to the theoretical maximum of 6.64. The "beyond suspicion" property is destroyed, however, because one destination is 100 times likelier than any other. Therefore, we also consider *min-entropy*, which captures the probability of the *likeliest* destination:

$$\mathsf{RAmin}_x = -\log_2(\max_{1 \leq j \leq H}(\mathsf{RAmsg}_x[j]))$$

3. CALCULATING ANONYMITY

Let $m_1, m_2, \ldots, m_M$ be the mix nodes that form the network. We assume that all routes have the same length L, and that mix nodes and destinations are distinct. (The model is easily adapted to other scenarios.) Recall that each sender S_i is sending messages $x_{i1}, \ldots, x_{in_i}$; let $N' = \sum_i n_i$. For each message x_i where $1 \leq i \leq N'$, its *flow* f_i is the time sequence of mix-to-mix messages that carry the contents of x_i through the network.

We assume that message are encrypted with a pairwise key when sent from mix to mix, and that the global attacker can observe the origin and destination of each (encrypted) message, yet this is *all* he can do. In particular, we ignore active attacks such as trickles and floods, intersection and long-term statistical attacks, traffic analysis, and so on. These attacks would result in even worse anonymity loss than shown by our experiments.

Consider the discretized timeline $t_1, \ldots, t_L$. Each "moment" corresponds to one hop in the messages' routes. At time t_j, a global attacker observes messages $q_{j1}, \ldots, q_{jN'}$ arriving to some subset of mixes. Each q_{ji} carries the contents of a single $x_{i'}$, and for each $x_{i'}$, there exists some q_{ji} that carries its contents. ($\forall i \; q_{1i} = x_i$.) This correspondence, however, is not directly observable by the attacker due to mixing performed at each hop.

Define the flow matrix $\mathsf{Flow} : N' \times L \rightarrow M$ so that $\mathsf{Flow}[i,j]$ contains the destination of message q_{ji}. Define $\mathsf{FlowCount}[l,j] = count_{1 \leq k \leq N'}(\mathsf{Flow}[k,j] = m_l)$ to be the number of messages entering mix m_l at time t_j and leaving it at time t_{j+1} (except for $j = L$, in which case all messages enter their final destinations).

Consider some message $x \in \{x_1, \ldots, x_{N'}\}$ that entered the network at time t_1. Let i_x be the index of this message. We recursively calculate the probability $\mathsf{Pmsg}_x[i,j]$ that message q_{ji} observed by the attacker at time t_j carries the contents of x. Initially, for all i

$$\mathsf{Pmsg}_x[i,1] = \begin{cases} 1 & \text{if } i = i_x \\ 0 & \text{otherwise} \end{cases}$$

Now consider time t_j such that $2 \leq j \leq L$, and messages q_{ji} where $1 \leq i \leq N'$. For every mix m_l where $1 \leq l \leq M$, define

$$\mathsf{PthruMix}_x[l,j] = \sum_{\forall i \in [1..N'] \text{ s.t.} \mathsf{Flow}[i,j] = m_l} \mathsf{Pmsg}_x[i,j]$$

This is the probability that one of the messages entering m_l at time t_j carries the contents of x. The attacker need not know *which* message this is. Then define

$$\mathsf{Pmsg}_x[i,j] = \begin{cases} \mathsf{Pmsg}_x[i,j-1] & \text{if } m_l \text{ is compromised} \\ \frac{\mathsf{PthruMix}_x[l,j-1]}{\mathsf{FlowCount}[l,j-1]} & \text{if } m_l \text{ is good} \end{cases}$$

Intuitively, this means that a good mix permutes its incoming messages and every outgoing message is equally likely to carry the contents of a given incoming message. By contrast, any permutation carried out by a compromised mix is completely visible to the attacker, and does not change the probability that a given message carries the contents of message x.

Finally, $\forall j, 1 \leq j \leq H$ where H is the number of destinations

$$\mathsf{RAmsg}_x[j] = \sum_{\forall k \in [1..N'] \text{ s.t. } \mathsf{Flow}[k,L] = D_j} \mathsf{Pmsg}_x[k,L]$$

This distribution allows us to compute relationship anonymity. Let $x_{i1}, \ldots, x_{in_i}$ be the messages sent by sender S_i at time t_1. The probability that S_i is communicating with destination D_j is

$$\mathsf{RA}_{ij} = 1 - \Pi_{1 \leq k \leq n_i}(1 - \mathsf{RAmsg}_{x_{ik}}[j])$$

4. EXPERIMENTS

We used the methodology described in section 3 to compute relationship anonymity in several types of mix networks. In our experiments, we assume that the number of destinations is equal to the number of senders. Not all destinations necessarily receive messages, while some destinations may receive multiple messages.

In each experiment, we randomly generate a 5-hop route for every message. We consider a *free-route* network, in which the destination of each hop is selected randomly from all mixes in the network, and a *stratified* network, inspired by [5, 6]. In the latter, the mixes are split into 5 groups, and the destination of the ith hop is selected randomly from the mixes in the ith group. For each route configuration, we calculate destination entropy and min-entropy for all senders, and average the results across 5,000 simulations.

When using maximum entropy for comparison purposes, we deliberately do *not* take non-uniform destination selection into account. Our goal is to demonstrate the effects of non-uniform destination selection by comparing anonymity when the distribution of destinations is skewed and when this distribution is uniform. The maximum entropy curve represents the latter on our plots.

When the actual destination distribution is not uniform, the maximum entropy associated with uniform destination distribution may not be achievable even by a perfect "black-box" anonymity system. This is precisely our point. An average user of a mix network may not know how other users select their destinations. In a truly peer-to-peer mix network, this information is not available even to operators of individual nodes. When a user knows only the total number of other users, network topology, and some assumption about the fraction of the mixes that are controlled by the adversary, anonymity under uniform destination selection provides a baseline against which actual anonymity can be compared.

Previous work on anonymity in large mix networks [6] assumed that all messages are distributed equally across network links at each hop. This is not realistic unless the ratio of messages to mixes is very large, which may not always be the case in practical scenarios. It is thus important to consider anonymity loss due to bad route selection. If the route selected by some user has few intersections with other users' routes, then this user will enjoy very poor anonymity even if all mixes on his route are trustworthy. As our experiments show, this effect is especially pronounced in networks where the number of users is comparable to the number of mixes, as will often be the case in peer-to-peer mix networks.

4.1 Impact of destination distribution

Unlike sender anonymity, relationship anonymity is sensitive to the distribution of messages' destinations. In our first set of experiments, we consider the uniformly random distribution (*i.e.*, each sender selects the destination uniformly at random from the fixed set of all available destinations) as well as Zipfian distributions with various skew parameters. In a Zipfian distribution, the relative probability of occurrence for the nth most popular element is $\frac{1}{n^{\alpha}}$, where α is the skew. It is believed that real-world Web browsing patterns are governed by a Zipfian distribution, *i.e.*, a small fraction of sites account for the majority of traffic destinations [2].

In fig. 1, we show the results for a 50-node free-route network where any node is compromised with a 25% probability. The number of senders increases from 50 (same as the number of mixes) to 500. In addition to measuring relationship anonymity, we also calculate sender anonymity for each message exiting the network.

Anonymity loss due to free route selection. Our first observation about fig. 1 is that there is a significant anonymity loss regardless of the destination distribution. This can be explained by statistical properties of random route selection in a free-route network.

For example, consider a free-route network with 50 senders and 50 mixes. Even if the destination for each message is selected uniformly at random (and relationship anonymity is thus close to simple sender anonymity for each message exiting the network), the effective anonymity set for the average user is 8 rather than 50. The primary reason for this is *not* that 25% of mixes are compromised. In order to mix well, routes of different messages must have many intersections, and, moreover, messages must arrive to the intersection mix at approximately the same time, which in our idealized model means the mix must occupy the same hop in their respective routes. When 50 routes of 5 nodes each are selected randomly from 50 mixes, the routes have relatively few "synchronized" intersections, resulting in very poor mixing. Quality of mixing, expressed as the ratio of actual anonymity to the theoretical maximum, improves as the number of senders and routes increases.

Anonymity loss due to skewed destination distribution. Our second observation is that skewed destination distributions are associated with very poor relationship anonymity. For example, a Zipfian distribution with the skew of 1 (in the case of 500 senders, this means roughly that 76% of senders are communicating with 20% of destinations) produces relationship anonymity entropy of 6 if there are 500 senders. This means that the effective size of the anonymity set is only 64, even though the size of the sender anonymity set for the average message exiting the network is close to 380 (which is still noticeably below 500 due to effects of route selection).

The explanation is very natural, but worth keeping in mind when using mix networks. If most users are communicating with a relatively small subset of destinations, then it is easier for the attacker to infer who a certain user is communicating with even if the network preserves sender anonymity of any given message. In other words, with highly skewed distributions typical of Web browsing patterns, *even a perfect mix network provides poor relationship anonymity*.

High entropy does not guarantee that the relationship is "beyond suspicion." Our third observation is that there is a substantial difference between anonymity measured as entropy and anonymity measured as min-entropy. This shows that in many situations entropy is not an adequate measure of anonymity. For example, consider the skew-1 Zipfian distribution in the case of 500 senders. Entropy of the distribution of potential destinations for the average message entering the network is around 6, which is not great, but still equivalent to an anonymity set whose effective size is 64.

Min-entropy, however, is less than 3 in this case, and the effective size of the anonymity set is less than 8. Very roughly, this indicates that the distribution of potential destinations has a lot of randomness, but some destinations are *significantly* likelier than others. Anonymity "beyond suspicion" is thus not achieved.

4.2 Impact of route selection algorithm

In our second set of experiments, we compare different route selection algorithms. It has been observed that increasing the number of feasible routes may result in worse anonymity [1]. Adding a mix to a route helps only if it receives multiple messages at the same time. When senders are limited to a few choices for each hop of their routes, multiple messages are more likely to go through the same mixes at the same hop, resulting in better anonymity.

In a stratified network, route selection is restricted to a subset of mixes at each hop, greatly reducing the number of feasible routes. It is thus to be expected that a stratified network provides better anonymity than a free-route network. (This was observed in [6].)

In fig. 2, we compute relationship anonymity for a free-route and a stratified network. The number of destinations is equal to the number of senders, and the senders' choice of destinations is governed by the Zipfian distribution with skew of 1.2.

When the number of senders is comparable to the number of mixes, there is a noticeable difference in anonymity between the free-route network and the stratified network. The stratified network provides better anonymity because the number of feasible routes is smaller. This difference gradually disappears as the number of routes increases, and anonymity loss due to non-uniform destination selection dominates the effects of route selection.

5. CONCLUSIONS

We presented a definition and calculation methodology for relationship anonymity in mix networks with arbitrary route configurations. As our simulations show, skewed destination distributions typical of Web browsing lead to a significant deterioration in relationship anonymity even when the number of compromised mixes is relatively small. Our experiments also demonstrate the importance of route selection algorithms. Unless the number of routes is much greater than the number of mixes, free route selection results in substantially worse anonymity than selection algorithms which restrict the number of feasible routes. Finally, our experiments show that high entropy of anonymity sets does not necessar-

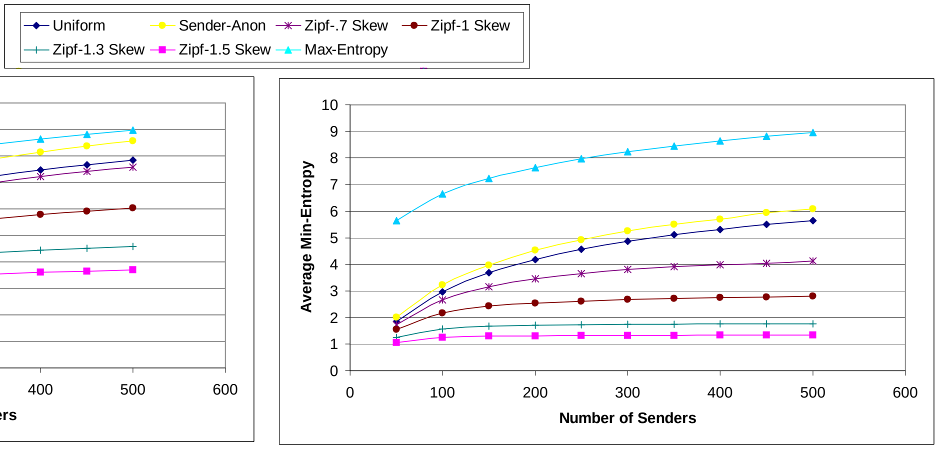

Figure 1: Relationship anonymity for different destination distributions (25% probability of node compromise).

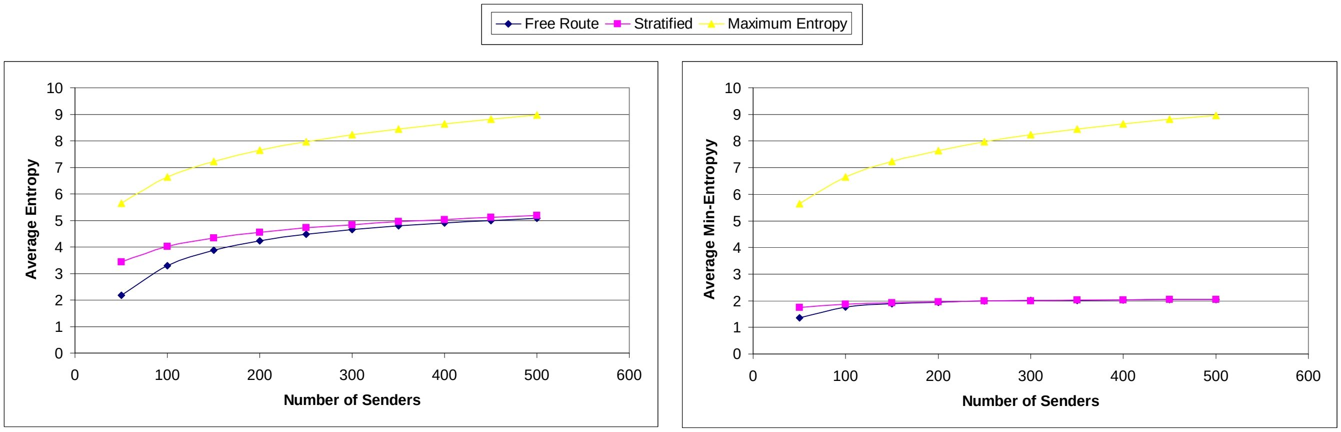

Figure 2: Relationship anonymity for different route selection algorithms (25% probability of node compromise).

ily mean that the network provides anonymity "beyond suspicion." Even in high-entropy networks, there can be significant differences between probabilities of different destinations. This calls for new anonymity metrics that better capture the "beyond suspicion" property in large anonymity sets. Finally, investigating susceptibility of different network topologies to long-term intersection attacks when multiple routes from the same sender to the same destination are established over time is an interesting topic of future research.

6. REFERENCES

[1] Berthold, O., Pfitzmann, A., and Standtke, R. The disadvantages of free MIX routes and how to overcome them. In *Proc. Workshop on Design Issues in Anonymity and Unobservability* (2000), vol. 2009 of *LNCS*, pp. 30–45.

[2] Breslau, L., Cao, P., Fan, L., Phillips, G., and Shenker, S. Web caching and Zipf-like distributions: evidence and implications. In *Proc. INFOCOM (Volume 1)* (1999), pp. 126–134.

[3] Chaum, D. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM 24*, 2 (1981), 84–88.

[4] Díaz, C., Seys, S., Claessens, J., and Preneel, B. Towards measuring anonymity. In *Proc. 2nd International Workshop on Privacy-Enhancing Technologies* (2002), vol. 2482 of *LNCS*, pp. 54–68.

[5] Dingledine, R., Freedman, M., Hopwood, D., and Molnar, D. A reputation system to increase MIX-net reliability. In *Proc. 4th International Workshop on Information Hiding* (2001), vol. 2137 of *LNCS*, pp. 126–141.

[6] Dingledine, R., Shmatikov, V., and Syverson, P. Synchronous batching: from cascades to free routes. In *Proc. 4th International Workshop on Privacy-Enhancing Technologies* (2004), vol. 3424 of *LNCS*, pp. 186–206.

[7] Pfitzmann, A., Köhntopp, M., and Shostack, A. Anonymity, unobservability, and pseudonymity - a proposal for terminology. Manuscript, June 2001.

[8] Reiter, M., and Rubin, A. Crowds: Anonymity for web transactions. *ACM Transactions on Information and System Security 1*, 1 (1998), 66–92.

[9] Serjantov, A., and Danezis, G. Towards an information theoretic metric for anonymity. In *Proc. 2nd International Workshop on Privacy-Enhancing Technologies* (2002), vol. 2482 of *LNCS*, pp. 41–53.

[10] Tóth, G., Hornák, Z., and Vajda, F. Measuring anonymity revisited. In *Proc. 9th Nordic Workshop on Secure IT Systems* (2004), pp. 85–90.